



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 0 981 228 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:  
23.02.2000 Bulletin 2000/08

(51) Int. Cl.<sup>7</sup>: H04L 12/56, H04Q 11/04

(21) Application number: 99109264.4

(22) Date of filing: 26.05.1999

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(30) Priority: 19.08.1998 CA 2245367

(71) Applicant:  
Newbridge Networks Corporation  
Kanata, Ontario K2K 2E6 (CA)

(72) Inventors:  
• Chow, Henry  
Kanata, Ontario K2K 2R7 (CA)

• Hung, Anthony  
Ottawa, Ontario K2B 5K8 (CA)  
• Janoska, Mark  
Carleton Place, Ontario K7C 3J1 (CA)  
• Ramaswamy, Srinivasan  
Nepean, Ontario K2B 2M8 (CA)

(74) Representative:  
Modiano, Guido, Dr.-Ing. et al  
Modiano, Josif, Pisanty & Staub,  
Baaderstrasse 3  
80469 München (DE)

(54) Two-component bandwidth scheduler having application in multi-class digital communication systems

(57) The method for servicing queues (14) holding data (12) packets for subsequent transmission to a communications link (16) processing comprises the steps of servicing each queue by forwarding its data packets to the link at time intervals corresponding to a guaranteed service rate of the queue, provided the queue is non-empty; and, during time intervals when none of the queues have packets being forwarded to the link in conformance with the above step, servicing the queues in accordance with a proportion of idle bandwidth allocated to each queue. The method is preferably carried out by a hierarchical scheduler (10) comprising an exhaustive scheduler (30) servicing a plurality of lower level schedulers in accordance with non-equal priority levels assigned thereto; a non-work conserving

shaper scheduler (20) feeding the exhaustive scheduler; and a work conserving idle bandwidth scheduler (25) feeding the exhaustive scheduler (30). Each queue (14) concurrently contends for service from the shaper scheduler (20) and the idle bandwidth schedulers (25). The shaper scheduler (20) servicing the queue has a higher priority level with respect to the exhaustive scheduler (30) than the idle bandwidth scheduler (25) servicing the same queue. The technique distributes the idle bandwidth of the communications link (16) in a way which is de-coupled from the guaranteed service rates of the queues, thereby providing a more efficient distribution of the total available bandwidth.

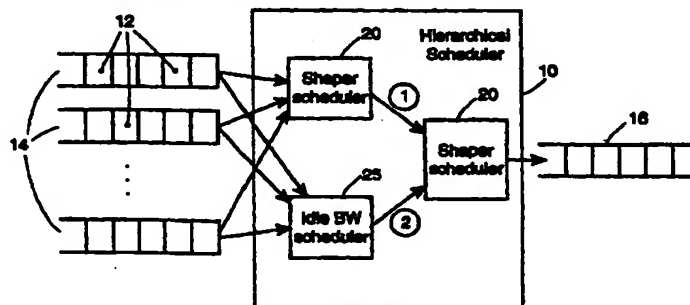


FIGURE 2

EP 0 981 228 A2

## Description

## Field of Invention

- 5 [0001] The invention generally relates to the art of scheduling systems wherein messages associated with plural processes are stored in a number of queues for subsequent processing by a single resource having limited processing capability. The invention has particular application to the field of digital communications systems and in this aspect relates to a scheduler and related method for efficiently allocating the bandwidth of a communications link amongst multiple queues which may be associated with a variety of service classes.

10

## Background of Invention

- [0002] In various types of communication systems, including Asynchronous Transfer Mode (ATM) systems, situations often arise where a number of connections vie for the bandwidth of a communications link in a communication device, such as at a network node. When such a situation arises, it is necessary to queue or buffer data packets or cells from the contending connections, and the queues must be serviced in some "fair" way in order to ensure that all of the connections are adequately serviced.

- [0003] A similar situation arises in the more general case where plural processes contend for a single resource. For instance, a distributed processing system may comprise a number of local controllers, responsible for various facets of the system, which are connected to a central controller, responsible for the overall management of the system. The local controllers communicate with the central controller by sending it messages, which the central controller must process, i.e., act upon. In this sense, the local controllers present "jobs" to the central controller. At any instant of time, some of the local controllers will not be busy, having no messages which must be processed by the central controller. Concurrently, some of the local controllers will be busy, presenting multiple messages, and hence potential jobs, to the central controller. Since the central controller may be busy with other jobs, it stores the messages in various queues, e.g., according to the type or class of local controller from which the message originated, until such time the central controller can process the message and carry out the associated job. These messages must also be serviced in some fair way to ensure that all of the local controllers are adequately handled. It will be seen from the foregoing that the messages or jobs correspond to data packets of the digital communication system, and the fixed processing power or speed of the central controller corresponds to the bandwidth of the communications link.

- [0004] A common "fair" scheduling scheme is proportional weighted fair queuing (hereinafter "proportional WFQ") wherein each queue, corresponding to each connection, is assigned a weight proportional to its allocated service rate. The proportional WFQ scheduler uses this weight to determine the amount of service given to the queue such that the scheduler is able to provide the allocated service rate for a given connection over a reasonably long busy period (i.e., when its queue is continuously non-empty), provided that the scheduler is not over-booked. The notion of an allocated service rate suits ATM systems in particular because almost all of the five currently defined ATM service classes rely on rate as a basis for defining quality of service (QoS). For instance, constant bit rate (CBR) connections are guaranteed a cell loss ratio (CLR) and delay for cells that conform to the peak cell rate (PCR). Variable bit rate (VBR) connections, real-time and non-real-time, are also guaranteed a CLR and delay for cells that conform to the sustained cell rate (SCR) and PCR. An available bit rate (ABR) connection is given a variable service rate that is between a minimum cell rate (MCR) and PCR. Unspecified bit rate (UBR) connections are associated with PCRs, and are soon anticipated to also be associated with MCRs.

- [0005] In addition to the allocated service rate, because a proportional WFQ scheduler is work conserving, each non-empty queue will also receive a certain amount of instantaneous idle bandwidth. This is the extra service bandwidth that a queue receives due to (1) any unallocated bandwidth of a communications link, and (2) any allocated but currently unused bandwidth arising from the idle, non-busy periods of the other queues at the contention point.

- [0006] To explain this in greater detail, suppose that queue  $n$  is given a weight  $\phi_n$  which is proportional to the allocated service rate queue  $n$  should receive. The proportional WFQ scheduler thus distributes the total allocated bandwidth of the communication link amongst all the queues in proportion to their allocated service rates. Consequently, the idle bandwidth of the link is also distributed in proportion to the allocated service rates of all the non-empty queues. An example of this is shown in Figure 1(a) where four queues 14, corresponding to four connections A, B, C & D, are serviced by a proportional WFQ multiplexer 8 in order to produce an output cell stream or link 16. Connections A, B & C have allocated service rates equal to 30% of the total bandwidth associated with the link 16 and are thus equally weighted. The allocated service rate of connection D is equal to 10% of the total bandwidth of link 16. Figure 1(b) is a bandwidth occupancy chart illustrating how the link bandwidth is allocated to the connections. From time  $t = 0$  to 8, each of the connections has cells requiring servicing and thus the instantaneous bandwidth received by each connection is 25% of the total bandwidth. At time  $t = 8$ , however, only connections B and D are non-empty having cells to be serviced, and thus the instantaneous idle bandwidth (now being 50% of the total bandwidth) is allocated to connections B & D in

proportion to their allocated service rates. Thus, at time  $t = 8$ , connection B receives 75% of the instantaneous total bandwidth and connection D receives 25% of the instantaneous total bandwidth. In general, the theoretical instantaneous service that queue  $n$  receives at time  $t$  when it is non-empty is  $\phi_n / \sum_{i \in A(t)} \phi_i$  where  $A(t)$  is the index set of non-empty queues at time  $t$ .

- 5 [0007] Suppose then that a proportional WFQ scheduler is used in an ATM communications device, such as a network node. A CBR connection should have an allocated service rate equal to its PCR. A VBR connection should have an allocated service rate, VBW (virtual bandwidth), which is at least equal to its SCR and less than its PCR. (VBW is typically statistically calculated at set up by the connection and admission control (CAC) function of a network.) An ABR connection should have an allocated service rate equal to its SCR, and a UBR connection should have an allocated service rate equal to zero. So, in such an scenario, the amount of idle bandwidth that a CBR connection sees is proportional to its PCR, and that an ABR connection sees is proportional to its MCR. This may result in very undesirable service. For example, suppose that a switch is carrying four connections (only): one is CBR with PCR = 980 kbps, two connections are ABR with MCR = 10 kbps, and one is UBR. Consequently, the idle bandwidth distribution is 98% for the CBR connection and 1% for each of the ABR connections, assuming a period when all the connections are busy.
- 10 Such a distribution is certainly not desirable, since CBR connections should generally not receive service bandwidth beyond their PCRs. ABR connections would get extra bandwidth in proportion to their MCRs; a phenomenon commonly termed MCR proportional service. MCR proportional service is one way of fairly distributing idle bandwidth fairly, but the literature has other methods such as MCR plus fair share which proportional WFQ cannot support. And the UBR connection only gets service if all the other queues are empty. Such distributions of the idle bandwidth are not optimal, and hence it is desired to achieve a more efficient distribution of the idle bandwidth.
- 20

#### Summary of Invention

- 25 [0008] Generally speaking, the invention provides a method for servicing a plurality of queues holding messages, such as data packets, destined for processing by a resource having a finite processing bandwidth, such as a communications link having a finite transmission bandwidth. The method comprises the steps of: (a) provisioning each queue with a minimum guaranteed service rate; (b) provisioning each queue with an idle bandwidth proportion; (c) servicing each queue by forwarding messages thereof to the resource at time intervals corresponding to the minimum guaranteed service rate of the queue, provided the queue is non-empty; and (d) servicing the queues in accordance with the proportion of idle bandwidth allocated to each queue during time intervals when none of the queues have packets being forwarded to the resource in conformance with step (c). In this manner, the amount of instantaneous idle bandwidth that a queue receives is decoupled from the allocated service rate granted to the queue.
- 30

- [0009] In the preferred embodiment, the above method is carried out by a hierarchical scheduler which comprises (a) an exhaustive scheduler servicing a plurality of lower level schedulers in accordance with non-equal priority levels assigned thereto; (b) a non-work conserving shaper scheduler feeding the exhaustive scheduler; and (c) a work conserving idle bandwidth scheduler feeding the exhaustive scheduler. The exhaustive scheduler is configured so that the shaper scheduler is given exhaustive priority over the idle bandwidth scheduler. The hierarchical scheduler is coupled to the queues such that each queue concurrently contends for service from the shaper scheduler and from the idle bandwidth scheduler.
- 35

- 40 [0010] The non-work conserving shaper scheduler, such as a virtual clock shaper described below, generates a stream of data packets at a constant average bit rate. Since the shaper scheduler servicing a particular queue (which may correspond to one connection) has a higher priority than the work conserving idle bandwidth scheduler, such as a WFQ scheduler, the queue is guaranteed its allocated service rate during its busy period. However, the shaper scheduler does not always submit messages (or in the preferred embodiment, the identity of queues) to the exhaustive scheduler because not all queues are busy at all times, and even if a queue is busy, it may not be eligible to be serviced due to the non-work conserving nature of shaping. These periods constitute the idle bandwidth of the resource. During this "idle" time, the lower priority work conserving idle bandwidth scheduler servicing the queue is able to feed the exhaustive scheduler. The idle bandwidth scheduler distributes this idle bandwidth in manner which is preferably non-dependent upon the guaranteed service rates allocated to the queues. In the preferred embodiment, the idle bandwidth scheduler partitions the instantaneous idle bandwidth in a fixed manner or ratio between QoS classes, and equally between all connections associated with a particular QoS class.
- 45
- 50

- [0011] In various embodiments described herein, the shaper scheduler and the idle bandwidth scheduler are also each preferably composed of a plurality of subschedulers in order to more flexibly accommodate the distribution of idle bandwidth in an ATM application environment, as explained in greater detail below.
- 55

- [0012] According to another broad aspect of the invention, there is provided a hierarchical scheduler for servicing a plurality of queues holding messages. This scheduler comprises an exhaustive sub-scheduler servicing a plurality of lower level sub-schedulers in accordance with non-equal priority levels assigned thereto; M non-work conserving shaper sub-schedulers feeding the exhaustive sub-scheduler; and N work conserving idle bandwidth sub-schedulers

feeding the exhaustive sub-scheduler. A given queue concurrently contends for service from one of the shaper sub-schedulers and from one of the idle bandwidth sub-schedulers, and the shaper subscheduler servicing the given queue has a higher priority level with respect to the exhaustive sub-scheduler than the idle bandwidth sub-scheduler servicing the given queue.

#### Brief Description of Drawings

[0013] The foregoing and other aspects of the invention will become more apparent from the following description of the preferred embodiments thereof and the accompanying drawings which illustrate, by way of example only, the principles of the invention. In the drawings;

Figure 1(a) is a diagram exemplifying the queue arbitration problem that a prior art proportional WFQ scheduler or multiplexer has to manage;

Figure 1(b) is a bandwidth occupancy chart showing how the link bandwidth is allocated to the queues over time by the prior art proportional WFQ scheduler under the conditions shown in Figure 1(a);

Figure 2 is a functional block diagram illustrating a hierarchical scheduler in accordance with a first preferred embodiment of the invention;

Figure 3 is a functional block diagram illustrating a hierarchical scheduler in accordance with a second preferred embodiment of the invention;

Figure 4 is a flowchart illustrating, at a high level, a method according to the first and second preferred embodiments for implementing the hierarchical scheduler shown in Figures 2 and 3;

Figure 5 is a flowchart illustrating a packet pre-processing stage of the flowchart of Figure 4 in greater detail;

Figure 6 is a flowchart illustrating an output processing stage of the flowchart of Figure 4 in greater detail in accordance with the first preferred embodiment;

Figure 7 is a flowchart illustrating an embellishment to the output processing stage shown in the flowchart of Figure 6 in accordance with the second preferred embodiment;

Figure 8 is a functional block diagram illustrating a hierarchical scheduler in accordance with a third preferred embodiment of the invention; and

Figure 9 is a functional block diagram illustrating a hierarchical scheduler in accordance with a fourth preferred embodiment of the invention.

#### Detailed Description of Preferred Embodiments

[0014] Figure 2 is a functional block diagram illustrating a hierarchical scheduler 10 in accordance with a first preferred embodiment. As described above, the task of scheduler 10 is to schedule data packets 12 stored in a plurality of input queues 14 to a limited resource, such as output communications link 16, which has a fixed bandwidth or service rate associated therewith. In the preferred embodiments, the data packets 12 are fixed length ATM or ATM-like cells. Furthermore, scheduler 10 is synchronous in that the data packets 12 are dequeued and transmitted to the communications link 16 at a rate corresponding to the fixed link bandwidth. In other words, the communications link 16 can be viewed as being logically divided into time slots, such that scheduler 10 dequeues one data packet 12 per time slot.

[0015] According to the first preferred embodiment, scheduler 10 comprises three "sub-schedulers" 20, 25 and 30, i.e., substantially independent schedulers, which are interconnected in a two-level, hierarchical, arrangement. (For notational purposes, sub-schedulers 20, 25 and 30 will be referred to merely as "schedulers" since these schedulers may themselves be composed of sub-schedulers, as shown, for example, in Figure 3.) An exhaustive scheduler 30 is disposed at the top or primary level which provides the output to the communications link 16. The primary exhaustive scheduler, which is work conserving, serves a secondary non-work conserving shaper scheduler 20, described in greater detail below, and a secondary work conserving idle bandwidth scheduler 25 (such as a WFQ scheduler) also described in greater detail below. The two secondary schedulers 20 and 25 concurrently service the queues 14. The exhaustive scheduler 30, also known in the art as a static or strict priority scheduler, services processes in accordance

with priority levels assigned thereto. Thus, at any given time, the primary exhaustive scheduler 30 services the secondary scheduler with the highest priority level, provided the latter is "busy", i.e., requires servicing. In the preferred embodiment, the secondary shaper scheduler 20 is assigned a higher priority level than the secondary idle bandwidth scheduler 25, thereby ensuring that the former will always be serviced ahead of the latter when busy.

5 [0016] The secondary schedulers 20 and 25 simultaneously "serve" the set of queues 14 by examining the queues and selecting one of them in order to dequeue a data packet therefrom. When the secondary schedulers 20 and 25 have selected a queue to serve, they do not remove data packets from the queues; instead queue identifiers are submitted to the primary scheduler 30. Thus, at each time slot, the secondary schedulers 25 and 30 simultaneously and independently submit the identifier of an eligible queue, if any, to the primary scheduler 30. Once the secondary sched-  
10 ulers 20 and 25 have submitted the queue identifiers, the primary scheduler 30 then serves the highest priority secondary scheduler which submitted an eligible queue by dequeuing the head-of-line (HOL) data packet from the queue identified by that secondary scheduler.

[0017] The secondary shaper scheduler 25 is preferably implemented as a virtual clock shaper similar to that described in Stilladios, D. and Varma, A., "A General Methodology for Designing Efficient Traffic Scheduling and Shap-  
15 ing Algorithms", *Proceedings of I.E.E.E. INFOCOM*, Japan, 1997 (hereinafter "Stilladios"), which is incorporated herein by reference. Each queue is provisioned or associated with an minimum guaranteed service rate, and the secondary shaper scheduler 25 provides a constant flow of data packets from each queue at its guaranteed service rate. Thus, if each queue 14 corresponds to a virtual connection, then each connection is provided with a constant average bit rate stream. The secondary shaper scheduler 25 is non-work conserving, and thus does not necessarily select a queue to  
20 be served every time slot.

[0018] The secondary idle bandwidth scheduler 25 is preferably implemented as a WFQ scheduler wherein each queue 14 is preferably provisioned or associated with a fixed WFQ weight corresponding to a pre-determined allocation of the instantaneous idle bandwidth of the communications link 16. Thus, for instance, each queue could be assigned with a WFQ weight of  $1/N$ , where  $N$  is the total number of queues at any time. The secondary idle bandwidth scheduler  
25 25 is work conserving, such that it submits an eligible queue to the primary scheduler 30 each and every time slot, provided that at least one queue has a packet stored therein.

[0019] It will be re-called that both secondary schedulers 20 and 25 simultaneously serve the queues 14. Since the shaper scheduler 20 has a higher priority than the work conserving scheduler 25, the queues are guaranteed their allocated service rates during their busy periods. However, at various times the shaper scheduler 20 does not always submit  
30 queues to the exhaustive scheduler 30 because 1) the shaper scheduler 20 is non-work conserving such that no queues require a packet to be submitted during a particular time period in order to maintain their minimum guaranteed service or shaper rates, or 2) all of the queues having non-zero shaper rates are empty during a particular time period (although queues having zero shaper rates may be busy). These time periods constitute the idle bandwidth of the communications link 16. During this "idle" time, the lower priority work conserving idle bandwidth scheduler 25 which always  
35 (provided not all of the queues are empty) selects an eligible queue every time slot is able to feed the primary scheduler 30, and hence distribute the idle bandwidth amongst the queues in accordance with the particular scheme provided by scheduler 25. In this way, the amount of instantaneous idle bandwidth a queue receives is decoupled from or independent of the guaranteed service rate assigned to it.

[0020] Other types of work conserving sub-schedulers may alternatively be used to allocate the idle bandwidth of a resource, such as communications link 16. However, generally speaking, since each queue contends for service from  
40 a shaper scheduler and from a work-conserving scheduler, wherein the former is granted exhaustive priority over the latter, then the work-conserving scheduler allocates a portion, or proportion, of the instantaneous idle bandwidth of the resource (hereinafter "idle bandwidth proportion") to the queue.

[0021] A preferred configuration for scheduler 10 for handling various ATM traffic service classes is shown in Table 1 below. In the preferred configuration, the idle bandwidth is "hard-partitioned" between the various ATM service classes, and each queue is associated with one virtual connection, which in turn is associated with one of the ATM traffic  
45 classes. The idle bandwidth proportion that the rtVBR, nrtVBR, ABR and UBR service classes receive are  $p_1$ ,  $p_2$ ,  $p_3$ , and  $p_4$ , respectively, where  $p_1 + p_2 + p_3 + p_4 = 1$ . The idle bandwidth allocated to each service class is divided equally amongst connections belonging to the same class. Alternatively, the idle bandwidth allocated to each service  
50 class could be divided amongst connections of the same class in proportion to their guaranteed minimum service rates. Other schemes are also possible.

Table 1

ATM Service Class	Guaranteed Minimum Service Rate per queue (shaper rate)	Idle Bandwidth Proportion per queue (WFQ weight)
CBR	PCR	$\infty$
rtVBR.1	VBW	$\infty$
rtVBR.2/3	VBW	$\rho_1/N_{rtVBR.2/3}$
nrVBR	VBW	$\rho_2/N_{nrVBR}$
ABR	MCR	$\rho_3/N_{ABR}$
UBR	MCR	$\rho_4/N_{UBR}$

[0022] The CBR and rtVBR.1 queues are assigned a weight of  $\infty$  to ensure that CBR and rtVBR.1 traffic receive priority service over non-real-time traffic and rtVBR.2/3 traffic. In practice, such sources are typically subject to usage parameter control (UPC), i.e., a type of policing function to ensure that a connection is abiding with its traffic contract, and thus tend not to be bursty (i.e., conforming to their CDVTs.) Alternatively, other means may be provided in a communications system for shaping such sources to ensure conformance to the traffic contract. See for example co-pending Canadian patent application no. 2,229,577, filed February 12, 1998, and assigned to the instant assignee. As such, CBR and nrVBR.1 connections will generally rarely utilize the instantaneous idle bandwidth managed by scheduler 25.

[0023] The rtVBR.2/3 service class is treated as a non-real time class with respect to the distribution of idle bandwidth because of the nature of its traffic contract which only requires an ATM communications device, such as a network node having connection admission control (CAC), to guarantee service to CLP 0 cells but not low priority CLP 1 cells. See, for instance, ATM FORUM doc. no. af-tm-0056.000, "Traffic Management Specification, Version 4.0", April 1996, all of which is incorporated herein by reference. CLP 0 cells are subject to two bucket UPC (i.e., PCR conformance and SCR conformance) and thus typically tend to be well-behaved. However, CLP 1 cells are only subject to one bucket UPC (i.e., only PCR conformance) such that an aggregate CLP0+1 stream is only constrained by a PCR bucket. Accordingly, it is possible for CLP 1 cells from numerous rtVBR.2/3 connections to arrive at the node in extended bursts. If the rtVBR.2/3 traffic class were granted an idle bandwidth proportion of  $\infty$  as the other real time traffic classes, then the CLP 1 cells could cause service starvation of the non-real time traffic such as UBR, thereby flooding the node. However, CLP 1 cells may in fact be discarded or alternatively delayed to the same (or worse) extent than that of the non-real time traffic. Thus, by assigning the rtVBR-2/3 traffic a finite idle bandwidth proportion similar to the non-real time traffic, this problem may be avoided.

[0024] Figure 3 shows a hierarchical scheduler 50 according to a second preferred embodiment. In this embodiment, the secondary shaper scheduler 20 is composed of two shaper sub-schedulers 20A and 20B, and the secondary idle bandwidth scheduler 25 is composed of two work conserving sub-schedulers 25A and 25B. Queues 14A, which are associated with the CBR and rtVBR.1 ATM traffic classes, are concurrently serviced by shaper sub-scheduler 20A and work conserving sub-scheduler 25A. Queues 14B, which are associated with the rtVBR.2/3 ATM traffic classes, are concurrently serviced by shaper sub-scheduler 20A and work conserving sub-scheduler 25B. Queues 14C, which are associated with the non-real time ATM traffic classes, are concurrently serviced by shaper sub-scheduler 20B and work conserving sub-scheduler 25B. The order of priority associated with the secondary sub-schedulers is 20A, 20B, 25A and 25B, from highest to lowest priority.

[0025] A preferred configuration for scheduler 50 for handling the various ATM traffic service classes is shown in Table 2 below. As before, each queue is associated with a single virtual connection and each queue/virtual connection is ensured its guaranteed minimum service rate, yet due to the priority of shaper subscheduler 20A over 20B, real time traffic is scheduled ahead of non-real time traffic when shaping. As before, the critical real time traffic (CBR and rtVBR.1) classes also receive idle bandwidth in preference to the other traffic classes. However, the second embodiment provides a mechanism for better managing the distribution of idle bandwidth amongst the critical real-time traffic. This is because the idle bandwidth that the CBR and rtVBR service classes receive relative to one another is  $\alpha_1 : \alpha_2$ , where  $\alpha_1 + \alpha_2 = 1$ , and the allocated idle bandwidth per service class is divided equally amongst connections belonging to the same class. In this manner, the idle bandwidth proportion allocated to a CBR or rtVBR.1 queue can be well controlled, unlike the first embodiment, while still giving these service classes preferential treatment in receiving idle bandwidth over the other service classes.



Table 2

ATM Service Class	Guaranteed Minimum Service Rate per queue (shaper rate)	Idle Bandwidth Proportion per queue (WFQ weight)
CBR	PCR	$\alpha_1 / N_{\text{CBR}}$
rtVBR.1	VBW	$\alpha_2 / N_{\text{rtVBR.1}}$
rtVBR.2/3	VBW	$\rho_1 / N_{\text{rtVBR.2/3}}$
nrtVBR	VBW	$\rho_2 / N_{\text{nrtVBR}}$
ABR	MCR	$\rho_3 / N_{\text{ABR}}$
UBR	MCR	$\rho_4 / N_{\text{UBR}}$

[0026] It will also be noted from the second embodiment that the rtVBR.2/3 queues are treated as a separate group, serviced by the high priority shaper subscheduler 20A along with the other real time classes, but serviced by the lowest priority idle bandwidth sub-scheduler 25B, similar to the non-real time traffic classes. This is done to prevent CLP 1 flooding for the reasons stated above.

[0027] Figure 4 is a flowchart illustrating, at a high level, a method according to the first and second preferred embodiments for implementing hierarchical schedulers 10 and 50. The preferred method uses a form of time stamping, as described in greater detail below, for keeping track of when queues should be serviced. Thus, in the event 60 of the arrival of a data packet 12, an input processing stage 62 issues the time stamp, as required. In parallel, an output processing stage 64 dequeues data packets 12 from queues 14 for transmission to the communications link 16, as well as issues time stamps under certain circumstances.

[0028] In the preferred method, each queue 14 is associated with two time stamps, called theoretical emission times (TETs). One of these time stamps (hereinafter "shaper TET") is used by the secondary shaper scheduler 20, and the other times stamp is used by the preferred secondary WFQ idle bandwidth scheduler 25 (hereinafter "WFQ TET"). The preferred method differs from the virtual clock shaping technique described in Stiliadios, *supra*, in that in the prior art reference, each packet has a time stamp associated with it whereas in the preferred embodiments time stamping is performed at the level of each queue. The inventors have found that time stamping per queue is likely more economical to implement in practice because of lower memory storage requirements.

[0029] Both time stamps of a queue change whenever a new packet reaches the head-of-line (HOL) position in the queue. This happens either when (a) a data packet arrives at an empty queue, or (b) a data packet has just been served and its queue has a following data packet waiting to be serviced which progresses to the HOL position.

[0030] If a packet has just been served and its queue has no more packets waiting, its time stamps are no longer valid; i.e., the queue is no longer considered eligible for scheduling by the secondary shaper scheduler 20 and the preferred WFQ secondary idle bandwidth scheduler 25 (including any sub-schedulers thereof).

[0031] According to the preferred method, when a packet arrives at an empty queue  $i$ , thereby placing the packet at the HOL position thereof, the shaper TET (i.e., for the purposes of the secondary shaper sub-schedulers) of a given queue  $i$  at a current time  $RTP$  (real time pointer) is:

$$TET_i = \max\{TET_i, RTP\} + \frac{1}{R_i} \quad (1)$$

where  $R_i$  is the shaper rate of queue  $i$ .

[0032] In the event a packet is dequeued in queue  $i$  such that another packet waiting in the queue reaches the HOL position, the shaper TET of queue  $i$  is:

$$TET_i = RTP + \frac{1}{R_i} \quad (2)$$

[0033] The shaper time stamp stays with the queue until it is served. At each time slot, the secondary shaper scheduler 20 (or shaper sub-schedulers 20A and 20B in the second embodiment) serves the queue with the smallest shaper

TET out of all eligible queues. A queue  $i$  is eligible if

$$TET_i - \frac{1}{R_i} \leq RTP. \quad (3)$$

[0034] Mathematically, the index,  $j$ , of the chosen queue is expressed as:

$$j = \arg \left\{ \min_i \left\{ TET_i \mid TET_i - \frac{1}{R_i} \leq RTP \right\} \right\}. \quad (4)$$

[0035] In the preferred embodiments, the WFQ sub-scheduler which functions as the secondary work conserving scheduler 25 (or sub-schedulers 25A and 25B thereof in the second embodiment) also time stamp each queue serviced. This can be accomplished with a self-clocked fair queuing scheme (SCFQ), described in greater detail, for instance, in Goyal et al., "Start-Time Fair Queuing: A Scheduling Algorithm for Integrated Services Packet Switching Networks", *IEEE/ACM Trans. Networking*, Vol. 5, No. 5, October 1970, and incorporated herein by reference (hereinafter "Goyal"). With SCFQ, when a new packet arrives at the HOL position of a queue  $i$ , the WFQ TET is:

$$TET_i = \max\{TET_i, VTP\} + \frac{1}{\phi_i}, \quad (5)$$

where  $VTP$  (virtual time pointer) is the WFQ TET value of the queue that was last served by the WFQ scheduler, and  $\phi_i$  is the WFQ weight of queue  $i$ . (Note that each WFQ scheduler has a  $VTP$  variable associated with it. In contrast, all shapers employ a common  $RTP$  variable since it is a measure of real time.)

[0036] Alternatively, the time stamping for the preferred WFQ sub-scheduler may be carried out with a start-time fair queuing scheme (SFQ) described in greater detail in Goyal, *supra*. With SFQ, when a new packet arrives at the HOL position of queue  $i$ , the WFQ TET is:

$$TET_i = \max\{TET_i + \frac{1}{\phi_i}, VTP\}, \quad (6)$$

where  $VTP$ , again, is the WFQ TET value of the queue that was last served, and  $\phi_i$  is the WFQ weight of queue  $i$ .

[0037] Figure 5 is a flowchart illustrating the input processing stage 62 (of Fig. 4) in greater detail. An initial step 70 enqueues the arriving packet in the appropriate queue 14. The incoming packets may be queued using a variety of queuing schemes, for instance, per priority queuing according to QoS or ATM traffic classes (such as shown in Fig. 3), per VC queuing, per port queuing, or combinations of the foregoing. Step 72 checks whether the packet enqueued in step 70 was placed at the head of the corresponding queue. If not, the process terminates. Otherwise, step 74 calculates the time stamp for the appropriate queue for subsequent use by the shaper scheduler 25 in accordance with equation (1). Similarly, step 76 calculates the time stamp for the appropriate queue for subsequent use by the preferred WFQ scheduler in accordance with one of equations (5) and (6), as desired.

[0038] Figure 6 is a flowchart illustrating the output processing stage 64 (of Fig. 4) in greater detail for one time slot. In an initial step 80, the secondary shaper and idle bandwidth schedulers 20 and 25 (which are preferably implemented as separate independent computer processes or threads of execution) concurrently and independently select a queue, if any, for servicing in the current time slot. In the case of the shaper scheduler 20, the queue selection is made in accordance with equation (4). The idle bandwidth scheduler 20 selects the queue having the lowest valued TET.

[0039] At step 82, each of the secondary shaper and idle bandwidth schedulers 20 and 25 concludes whether or not a packet should be dequeued. If so, then at step 84, the secondary schedulers 20 and 25 submit the identities,  $Q_j$  and  $Q_k$ , of the respectively selected queues to the exhaustive scheduler 30.

[0040] At step 86, the primary exhaustive scheduler 30 checks whether or not the secondary shaper scheduler 20 has submitted a queue for servicing. If so, then at step 90 the primary exhaustive scheduler 30 dequeues the HOL packet from queue  $Q_j$  and forwards the packet to the communications link 16. At step 96, queue  $Q_j$  (the queue just served) is examined to determine whether it is empty or not. If it is not empty, then a following packet within queue  $Q_j$  becomes an HOL packet and thus at step 100 the shaper TET for queue  $Q_j$  is updated in accordance with equation (2).



However if queue  $Q_j$  is empty, then control passes to step 104. In this event an adjustment is made to the WFQ TET for queue  $Q_j$  as follows:

$$TET_{Q_j} = TET_{Q_j} - \frac{1}{\phi_{Q_j}}$$

This adjustment is made because time stamping system operates on the principle of accumulated "credits", as will be understood to those skilled in the art, and thus since the secondary idle bandwidth scheduler 20 did not in fact service queue  $Q_j$ , its credit in respect of the secondary scheduler 20 is preferably corrected before the next packet, which will be an HOL packet, arrives in queue  $Q_j$ .

[0041] If at step 86 the secondary shaper scheduler 20 has not submitted a queue for servicing, then at step 88 the primary exhaustive scheduler 30 checks whether or not the secondary idle bandwidth scheduler 25 has submitted a queue for servicing. If so, then at step 92 the primary exhaustive scheduler 30 dequeues the HOL packet from queue  $Q_K$  and forwards the packet to communications link 16. At step 98, queue  $Q_K$  (the queue just served) is examined to determine whether it is empty or not. If it is not empty, then a following packet within queue  $Q_K$  becomes an HOL packet and thus at step 102 the WFQ TET time stamp for queue  $Q_K$  is updated in accordance with one of equations (5) and (6) as desired. However if queue  $Q_K$  is empty, then control passes to step 106. In this event, a credit adjustment for the reasons previously described is now made to the shaper TET as follows:

$$TET_{Q_K} = TET_{Q_K} - \frac{1}{R_{Q_K}}$$

[0042] Figure 7 shows the processing carried out by the primary exhaustive scheduler 30 specifically for the case of the second embodiment (Fig. 4). It will be noted that a query is made of each of the secondary sub-schedulers 20A, 20B, 25A, and 25B, in that order, to determine whether or not they have submitted queues, and are serviced accordingly in a manner similar to that described above.

[0043] Figure 8 shows a tertiary level hierarchical scheduler 110 according to a third preferred embodiment. In this embodiment, the scheduler 110 services sixty-four inputs  $I_1 - I_{64}$ , each of which carries multiple connections from any of the various ATM traffic classes. The cells from each input are stored in a set of queues. Each set comprises at least one queue from each traffic class. Thus, for example, there are sixty-four CBR queues in total, one for each input; and sixty-four nrtVBR queues in total, one for each input. UBR traffic is divided into two sections: "UBR m" refers to multicast UBR traffic; "UBR s" refers to single-cast UBR traffic.

[0044] In this embodiment, the shaper scheduler 20 is composed of two shaper sub-schedulers 20A and 20B which feed an exhaustive sub-scheduler 112 (which in turn feeds exhaustive scheduler 30). Shaper sub-scheduler 20A services the real time traffic classes, and shaper sub-scheduler 20B services the non-real time traffic classes (and rtVBR.2/3). Shaper sub-scheduler 20A has a higher priority with respect to exhaustive sub-scheduler 112 than shaper sub-scheduler 20B, in a manner similar to the second preferred embodiment shown in Figure 3.

[0045] The idle bandwidth scheduler 25 is composed of five tertiary WFQ sub-schedulers 114A-E which feed a secondary WFQ sub-scheduler 116. In this hierarchy, the weights of the secondary WFQ sub-scheduler 116 correspond to a portion of idle bandwidth allocated to an ATM traffic group as a whole, and the weights of a given tertiary WFQ sub-scheduler 114 divide up the portion of idle bandwidth allocated to the corresponding ATM traffic group amongst the queues of that group.

[0046] A preferred configuration for scheduler 110 is shown in Table 3, below. In the table, the notation " $\Sigma_{rtVBR.2/3, \text{queue}} \text{VBW}$ " means, for example, the sum of the VBWs of all connections sharing a particular rtVBR.2/3 queue, and " $\Sigma_{rtVBR.2/3} \text{VBW}$ " means the sum of the VBWs of all rtVBR.2/3 connections on all inputs, i.e., in the entire system. In addition, the notation " $N_{UBR, \text{queue}}$ " and " $N_{UBR}$ " respectively denote, for example, the number of UBR connections in a particular queue, and the total number of UBR connections on all inputs.

Table 3

ATM Service Class	Guaranteed Minimum Service Rate per queue (shaper rate)	Tertiary WFQ weight (per queue)	Secondary WFQ Weight (per class)
CBR	$\Sigma_{\text{CBR,queue}} \text{PCR}$	$\Sigma_{\text{CBR,queue}} \text{PCR} / \Sigma_{\text{CBR}} \text{PCR}$	$\infty$
rtVBR.1	$\Sigma_{\text{rtVBR.1,queue}} \text{VBW}$	$\Sigma_{\text{rtVBR.1,queue}} \text{VBW} / \Sigma_{\text{rtVBR.1}} \text{VBW}$	
rtVBR.2/3	$\Sigma_{\text{rtVBR.2/3,queue}} \text{VBW}$	$\Sigma_{\text{rtVBR.2/3,queue}} \text{VBW} / \Sigma_{\text{rtVBR.2/3}} \text{VBW}$	$P_1$
nrtVBR	$\Sigma_{\text{nrtVBR,queue}} \text{VBW}$	$\Sigma_{\text{nrtVBR,queue}} \text{VBW} / \Sigma_{\text{nrtVBR}} \text{VBW}$	$P_2$
ABR	$\Sigma_{\text{ABR,queue}} \text{MCR}$	$N_{\text{ABR,queue}} / N_{\text{ABR}}$	$P_3$
UBR	$\Sigma_{\text{UBR,queue}} \text{MCR}$	$N_{\text{UBR,queue}} / N_{\text{UBR}}$	$P_4$

[0047] As before, the ATM service classes are divided into three distinct groups as follows: (1) critical real time traffic (CBR and rtVBR.1 queues); (2) real-time traffic with CLP1 cells that can be dropped (rtVBR.2 and rtVBR.3), and (3) non-real time traffic (nrtVBR, UBR and ABR.) The first and second groups (queues associated with real time traffic) have exhaustive priority over the third group (queues associated with non-real time traffic) with respect to shaper sub-schedulers 20A and 20B. This ensures that the real-time queues get the bandwidth allocated to them by the CAC with minimal delay.

[0048] With respect to the allocation of the idle bandwidth, the first group of critical real time traffic queues (whose connections have QoS guarantees for both CLP 0 and CLP 1 cells) receive exhaustive priority over the non-real time traffic since they are assigned a weight of  $\infty$ . The rtVBR.2/3 traffic class is not included in this category as rtVBR.2/3 traffic has CLP 1 cells which have no QoS guarantees. Instead, the idle bandwidth is distributed in a finite manner by the secondary WFQ sub-scheduler 116 by giving the aggregate rtVBR.2/3, nrtVBR, UBR and ABR traffic some predetermined ratio,  $p_i$ , of the total available idle bandwidth. Within the rtVBR.2/3 and nrtVBR traffic classes, the idle bandwidth of the service class is divided amongst the queues in proportion to the guaranteed service rates of the connections carried by the queues. Within the ABR and CBR classes, the idle bandwidth allocated to each class is distributed approximately equally amongst all connections in the class. Of course, it is also possible to achieve proportional distribution of the idle bandwidth by setting the weights of the tertiary sub-schedulers 114D and 114E for each queue to  $\Sigma_{\text{nrtVBR,queue}} \text{MCR} / \Sigma_{\text{nrtVBR}} \text{MCR}$ .

[0049] The two-level hierarchy of scheduler 25 also provides a more flexible arrangement for distributing the idle bandwidth allotment within a given service class. For example, consider a scenario where there are five queues in total in the nrtVBR class, and the class is allocated 20% of the idle bandwidth. If all five queues are active, then each queue receives  $1/5^{\text{th}}$  of 20% of the allocated idle bandwidth. However, if only one of the five queues is active, then due to the work conserving nature of WFQ, that queue would always be served by the appropriate tertiary WFQ sub-scheduler 114 which, in turn, receives 20% of the idle bandwidth from the secondary WFQ sub-scheduler 116. Thus, the only busy queue in the class receives 20% of the idle bandwidth. In contrast, in the second embodiment shown in Figure 3, each queue has an idle bandwidth proportion of  $p/N$ , and thus under the presented scenario the busy queue would only receive approximately  $1/5^{\text{th}}$  of 20% of the idle bandwidth (with the other 16% divided up amongst all busy schedulers).

[0050] The preferred hierarchical schedulers described above have been configured relative to an ATM applications environment. It will be understood that the hierarchical schedulers can be configured for other types applications. In the general case, the hierarchical schedulers may comprise M non-work conserving shaper sub-schedulers feeding an exhaustive sub-scheduler, and N work conserving idle bandwidth sub-schedulers feeding the exhaustive sub-scheduler, provided that a queue is concurrently contending for service from one of the shaper sub-schedulers and from one of the idle bandwidth sub-schedulers, and the shaper sub-scheduler serving the queue has a higher exhaustive priority level than the idle bandwidth sub-scheduler serving the same queue. Figure 9 shows an example of such a hierarchical scheduler where  $M = 2$  and  $N = 3$ .

[0051] It should also be understood that the communications link 16 may represent a transmission media, or alternatively, an internal bus or pathway in a communications device. Similarly, it should be appreciated that the preferred embodiments may be applied in other types of applications where a number of processes generate various types of messages or jobs which are stored in a plurality queues for subsequent processing by a single resource having a finite processing capability or bandwidth. Moreover, while the preferred embodiments have made reference to synchronous sub-schedulers operating on fixed length data packets over equal time periods, in alternative embodiments of the invention the schedulers may be asynchronous in nature and operate on variable length data packets. Similarly,

those skilled in the art will appreciate that other modifications and variations may be made to the preferred embodiments disclosed herein without departing from the spirit of the invention.

[0052] Where technical features mentioned in any claim are followed by reference signs, those reference signs have been included for the sole purpose of increasing the intelligibility of the claims and accordingly, such reference signs do not have any limiting effect on the scope of each element identified by way of example by such reference signs.

# Claims

1. A method for servicing a plurality of queues holding data packets destined for transmission over a communications link, said method comprising:

- (a) provisioning each queue with a minimum guaranteed service rate;
- (b) provisioning each queue with an idle bandwidth proportion;
- (c) servicing each queue by forwarding the data packets thereof to the communications link at time intervals approximately equivalent to the minimum guaranteed service rate of the queue, provided the queue is non-empty; and
- (d) during time intervals when none of the queues have messages being forwarded to the resource in conformance with step (c), servicing the queues in accordance with the proportion of idle bandwidth allocated to each queue.

2. A method according to Claim 1, wherein in step 1(d) the queues are serviced using a work conserving, weighted fair queuing (WFQ) scheme, and the idle bandwidth proportion is a WFQ weight.

3. A method according to claim 2, wherein the idle bandwidth WFQ weight for each queue is one of:

- (a) a predetermined fixed ratio relative to other queues;
- (b) proportional to the instantaneous length of the queue;
- (c) equal to  $P/N$ , where  $N$  is the number of queues in a common service class and  $P$  is a percentage of idle bandwidth allocated to that service class; and
- (d) proportional to its minimum guaranteed service rate relative to the guaranteed service rates of other queues within a common service class.

4. A method according to Claim 2 or Claim 3 wherein the data packets are of fixed size and the communications link has a fixed bandwidth rate.

5. Apparatus for servicing a plurality of queues holding data packets destined for transmission over a communications link, said apparatus comprising:

means for provisioning each queue with a minimum guaranteed service rate;

means for provisioning each queue with an idle bandwidth proportion;

first means for servicing each queue by forwarding the data packets thereof to the communications link at time intervals substantially corresponding to the minimum guaranteed service rate of the queue, provided the queue is non-empty; and

second means for servicing the queues in accordance with the proportion of idle bandwidth allocated to each queue during time intervals when none of the queues have messages being forwarded to the resource by the first means for servicing the queues.

6. Apparatus according to claim 5, wherein the first and second service means include a hierarchical scheduler comprising:

an exhaustive priority scheduler for servicing a plurality of lower level schedulers in accordance with non-equal priority levels assigned thereto;

a non-work conserving shaper scheduler feeding the exhaustive scheduler; and

a work conserving idle bandwidth scheduler feeding the exhaustive scheduler;

wherein a given queue is concurrently serviced by the shaper scheduler and the idle bandwidth scheduler, and wherein the shaper scheduler has a higher priority level with respect to the exhaustive sub-scheduler than the idle bandwidth scheduler.

7. Apparatus according to Claim 5, wherein the first and second service means include a hierarchical scheduler comprising:

an exhaustive scheduler servicing a plurality of lower level schedulers in accordance with non-equal priority levels assigned thereto;

M non-work conserving shaper schedulers feeding the exhaustive scheduler; and

N work conserving idle bandwidth schedulers feeding the exhaustive scheduler;

wherein a given queue is concurrently serviced by one of the shaper schedulers and one of the idle bandwidth schedulers, and wherein the shaper scheduler servicing the given queue has a higher priority level with respect to the exhaustive scheduler than the idle bandwidth scheduler servicing the given queue.

8. Apparatus according to Claim 6 or Claim 7, wherein a given shaper scheduler is itself hierarchical, comprising a multi-tiered arrangement of subschedulers, and a given idle bandwidth scheduler is itself hierarchical, comprising a multi-tiered arrangement of sub-schedulers, and wherein a set of queues are concurrently examined by one of the shaper sub-schedulers and one of the idle bandwidth sub-schedulers, and wherein the shaper sub-scheduler examining the set of queues has a higher priority level with respect to the exhaustive sub-scheduler than the idle bandwidth sub-scheduler examining the same set of queues.

9. Apparatus according to any of Claims 6, 7 and 8, wherein the schedulers service their respective queues by concluding which queue should be serviced and passing a queue identifier to a higher level scheduler, provided it exists.

10. Apparatus according to any of Claims 6 - 9, wherein the idle bandwidth scheduler is a weighted fair queue (WFQ) scheduler.

11. Apparatus according to Claim 10, wherein the queues are time-stamped and the time stamp, TET, of a given queue  $i$ , in respect of a shaper scheduler is set to

(a)

$$TET_i = \max\{TET_i, RTP\} + \frac{1}{R_i},$$

in the event an incoming data packet is placed at a head-of-line position in queue  $i$ ; or

(b)

$$TET_i = RTP + \frac{1}{R_i},$$

In the event a data packet is dequeued placing another packet at a head-of-line position in queue  $i$ , where RTP corresponds to a current time and  $R_i$  is the minimum guaranteed service rate of queue  $i$ ; and the shaper scheduler selects a queue,  $j$ , for service such that

$$j = \arg \left\{ \min_i \left\{ TET_i \mid TET_i - \frac{1}{R_i} \leq RTP \right\} \right\}.$$

12. Apparatus according to Claim 10, wherein the queues are time-stamped and the time stamp, TET, of a given queue  $i$ , in respect of a WFQ scheduler is set to one of (a)

$$TET_i = \max\{TET_i, VTP\} + \frac{1}{\phi_i}$$

5 and  
(b)

$$TET_i = \max\{TET_i + \frac{1}{\phi_i}, VTP\},$$

10

in the event an incoming data packet is placed at a head-of-line position in queue  $i$ , where  $VTP$  is the TET value of a queue that was last served, and  $\phi_i$  is a WFQ weight; and wherein the WFQ scheduler selects a queue,  $j$ , for service such that

15

$$j = \arg\{\min_i \{TET_i\}\}.$$

20

13. Apparatus according to Claims 11 or Claim 12, wherein: (a) in the event a packet is dequeued from a queue,  $Q_j$ , selected by the shaper scheduler, the TET for queue  $Q_j$  for the corresponding WFQ scheduler is adjusted by

25

$$TET_{Q_j} = TET_{Q_j} - \frac{1}{\phi_{Q_j}};$$

30

and (b) in the event a message is dequeued from a queue,  $Q_k$ , selected by the WFQ scheduler, the TET for queue  $Q_k$  for the corresponding shaper scheduler is adjusted by

35

$$TET_{Q_k} = TET_{Q_k} - \frac{1}{R_{Q_k}}.$$

14. Apparatus according to any of Claims 10 - 13, wherein the packets are of fixed size and the first and second means for servicing the queues operate synchronously to deliver one packet from the queues to the communications link during one fixed period time slot.

40

15. Apparatus according to any of Claims 10- 15, wherein the WFQ weight for each queue is one of:

45

- (a) a predetermined fixed ratio relative to other queues;
- (b) proportional to the instantaneous length of the queue;
- (c) equal to  $P/N$ , where  $N$  is the number of queues in a common service class and  $P$  is a percentage of idle bandwidth allocated to that service class; and
- (d) proportional to its minimum guaranteed service rate relative to the guaranteed service rates of other queues within a common service class.

50

16. Apparatus according to Claim 7, wherein the communications link is an ATM link;  $M=2$  and  $N=2$ ; and wherein queues associated with CBR and rtVBR.1 ATM traffic classes are serviced by a first shaper scheduler and a first WFQ idle bandwidth scheduler, queues associated with rtVBR.2/3 ATM traffic classes are serviced by the first shaper scheduler and a second WFQ idle bandwidth scheduler, and queues associated with nrtVBR.2/3, ABR, and UBR ATM traffic classes are serviced by a second shaper scheduler and the second WFQ idle bandwidth scheduler, and wherein the first shaper scheduler has exhaustive priority over the second shaper scheduler which has exhaustive priority over the first WFQ idle bandwidth scheduler which has exhaustive priority over the second WFQ idle bandwidth scheduler.

55

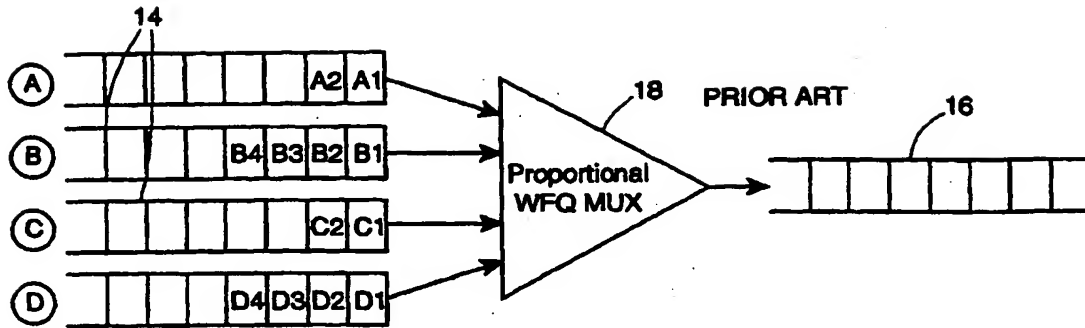


FIGURE 1(a)




FIGURE 1(b)






FIGURE 2




FIGURE 3




FIGURE 4

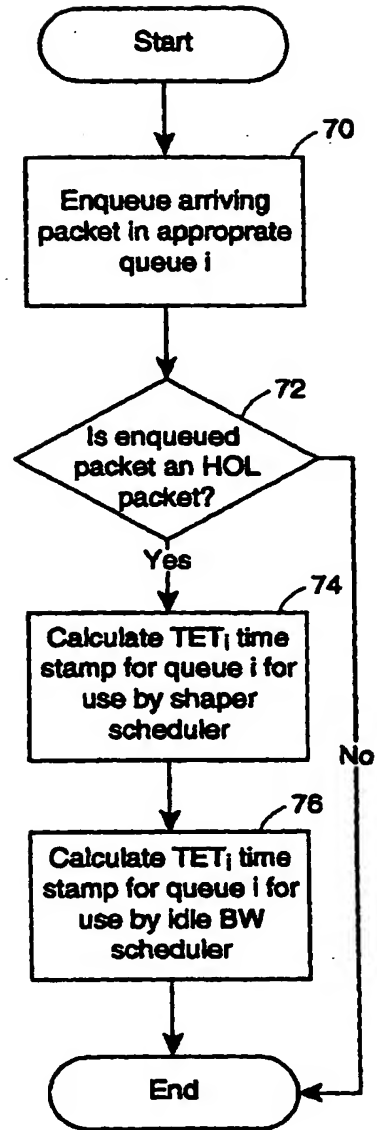


FIGURE 5

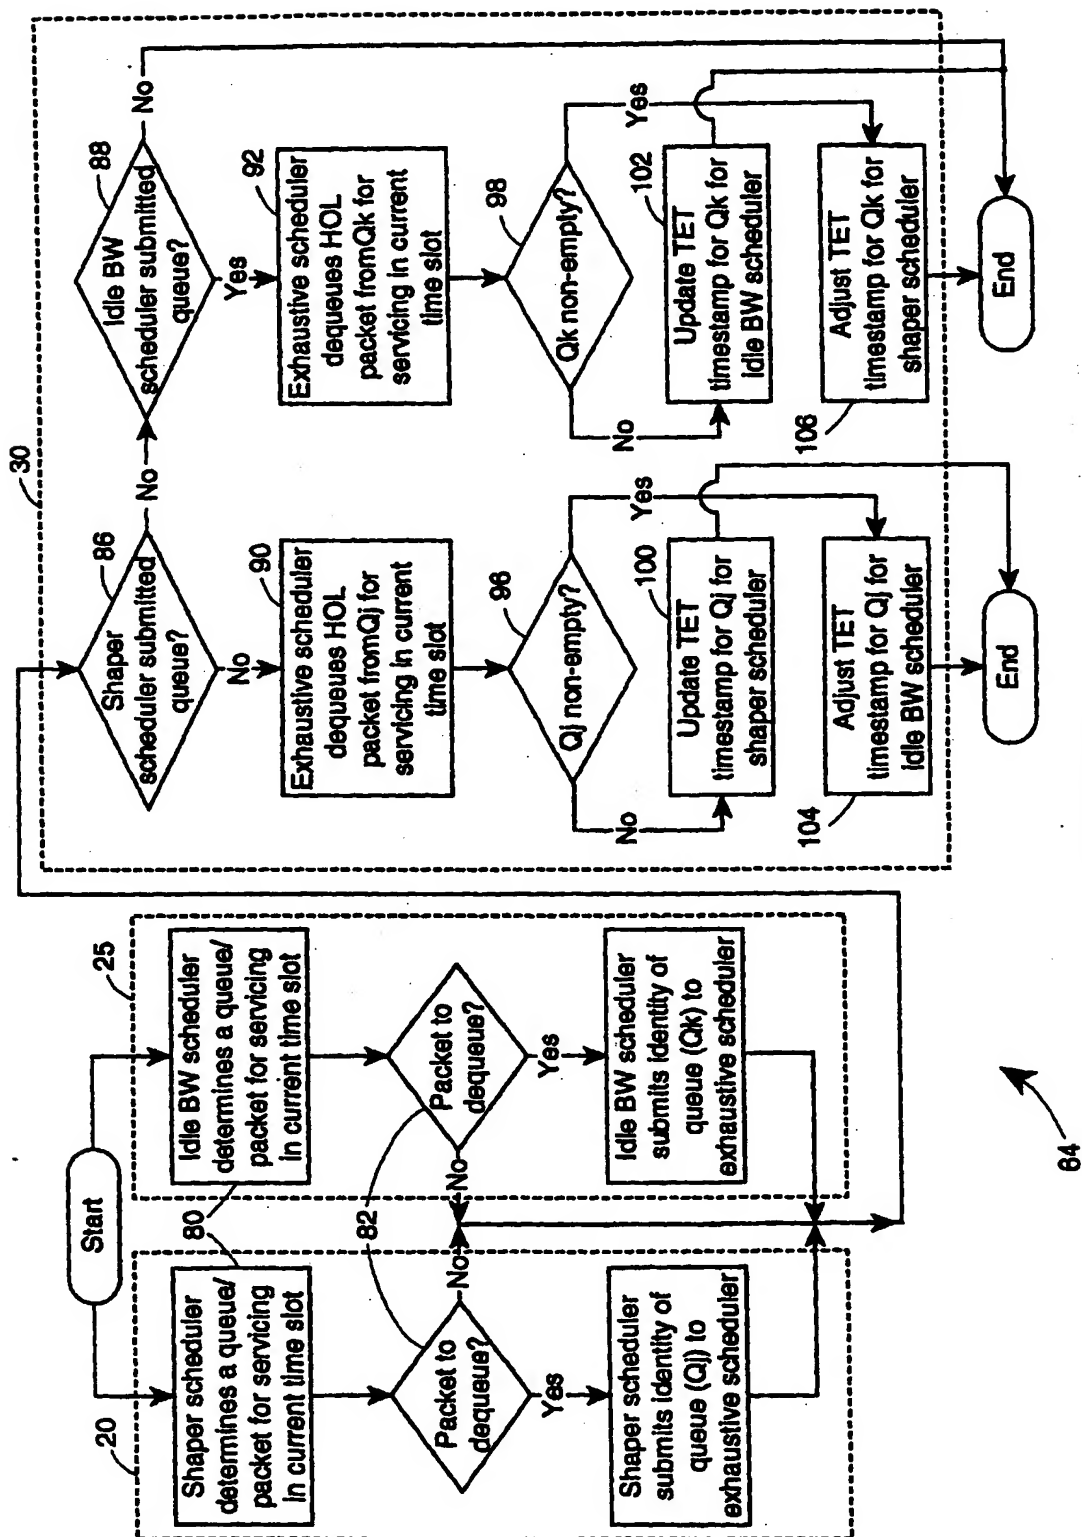


FIGURE 6

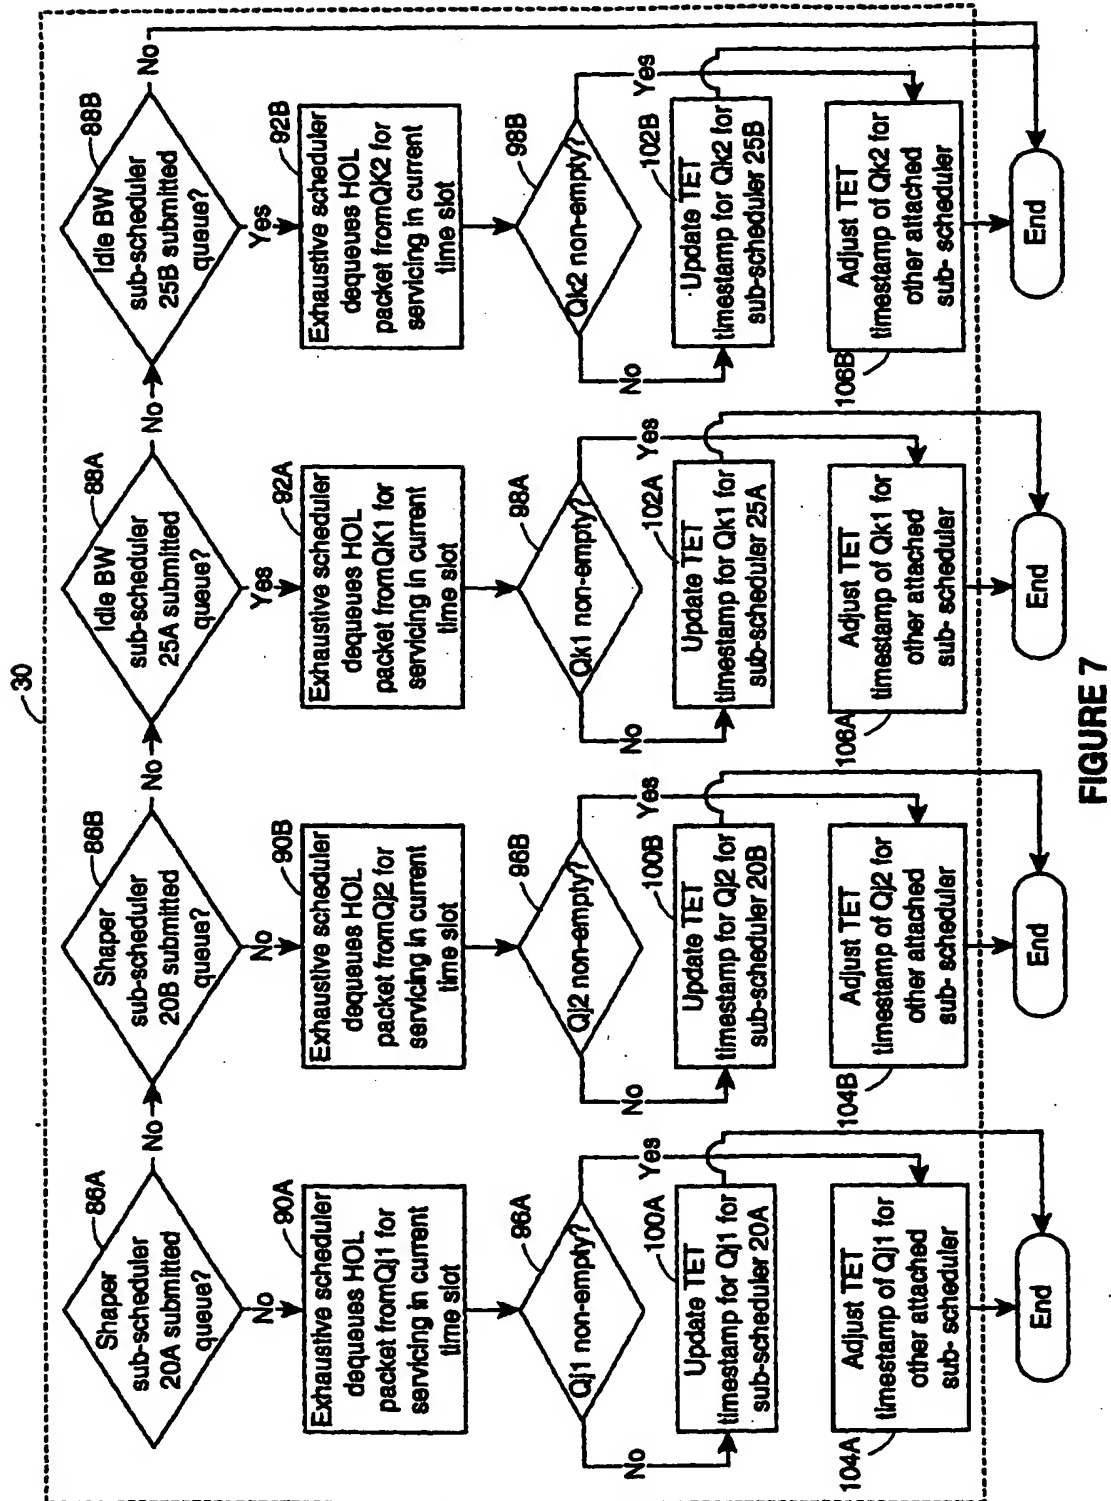


FIGURE 7




FIGURE 8

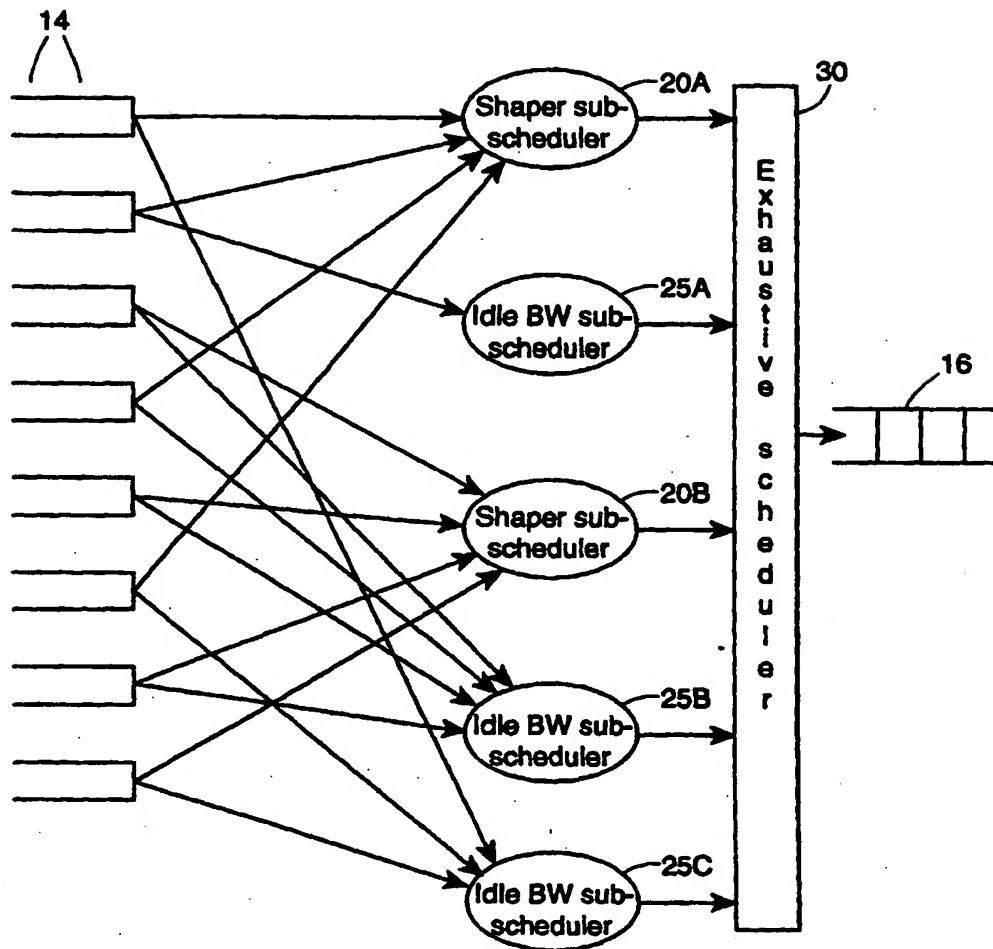


FIGURE 9